

Bayesian Inversion for Nonlinear Imaging Models Using Deep Generative Priors

Pakshal Bohra , Thanh-an Pham , Jonathan Dong , *Member, IEEE*, and Michael Unser , *Fellow, IEEE*

Abstract—Most modern imaging systems incorporate a computational pipeline to infer the image of interest from acquired measurements. The Bayesian approach to solve such ill-posed inverse problems involves the characterization of the posterior distribution of the image. It depends on the model of the imaging system and on prior knowledge on the image of interest. In this work, we present a Bayesian reconstruction framework for nonlinear imaging models where we specify the prior knowledge on the image through a deep generative model. We develop a tractable posterior-sampling scheme based on the Metropolis-adjusted Langevin algorithm for the class of nonlinear inverse problems where the forward model has a neural-network-like structure. This class includes most practical imaging modalities. We introduce the notion of augmented deep generative priors in order to suitably handle the recovery of quantitative images. We illustrate the advantages of our framework by applying it to two nonlinear imaging modalities—phase retrieval and optical diffraction tomography.

Index Terms—Bayesian inference, nonlinear inverse problems, phase retrieval, optical diffraction tomography, deep learning, neural networks, generative models, generative adversarial networks.

I. INTRODUCTION

IN PRACTICAL imaging systems, the object of interest $\mathbf{s} \in \mathbb{R}^K$ is observed indirectly by performing a series of measurements $\mathbf{y} \in \mathbb{C}^M$. Mathematically, this process is often modeled as

$$\mathbf{y} = \mathbf{H}(\mathbf{s}) + \mathbf{n}, \quad (1)$$

where $\mathbf{H} : \mathbb{R}^K \rightarrow \mathbb{C}^M$ is an operator that describes the physics of the imaging system and $\mathbf{n} \in \mathbb{C}^M$ is an additive noise. The operator \mathbf{H} can be linear or nonlinear, depending on the imaging

Manuscript received 5 July 2022; revised 17 November 2022; accepted 24 December 2022. Date of publication 11 January 2023; date of current version 23 January 2023. This work was supported in part by Swiss National Science Foundation under Grant 200020_184646/1 and in part by European Research Council (ERC Project FunLearn) under Grant 101020573. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Mauricio Delbracio. (*Corresponding author: Pakshal Bohra.*)

Pakshal Bohra, Jonathan Dong, and Michael Unser are with the Biomedical Imaging Group, École polytechnique fédérale de Lausanne, 1015 Lausanne, Switzerland (e-mail: pakshal.bohra@epfl.ch; jonathan.dong@epfl.ch; michael.unser@epfl.ch).

Thanh-an Pham was with the Biomedical Imaging Group, École polytechnique fédérale de Lausanne, 1015 Lausanne, Switzerland. He is now with the 3D Optical Systems Group, Department of Mechanical Engineering, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: tam-pham@mit.edu).

This article has supplementary downloadable material available at <https://doi.org/10.1109/TCI.2023.3236155>, provided by the authors.

Digital Object Identifier 10.1109/TCI.2023.3236155

modality. For example, in magnetic resonance imaging, one captures noisy samples of the Fourier transform of the signal. The task at hand is then to reconstruct the signal \mathbf{s} from the obtained measurements \mathbf{y} . Typically, such inverse problems are ill-posed, in the sense that there exist a multitude of signals which produce identical measurements. Thus, one cannot rely on direct inversion techniques to obtain relevant solutions.

A. Variational Methods

In variational methods, the solution to the inverse problem is specified as the minimizer of a cost functional

$$\hat{\mathbf{s}} = \arg \min_{\mathbf{s} \in \mathbb{R}^K} (E(\mathbf{y}, \mathbf{H}(\mathbf{s})) + \tau R(\mathbf{s})), \quad (2)$$

where the data-fidelity term $E : \mathbb{C}^M \times \mathbb{C}^M \rightarrow \mathbb{R}_+$ forces the solution to be consistent with the measurements, the regularization $R : \mathbb{R}^K \rightarrow \mathbb{R}_+$ imposes some prior constraints on the solution, and $\tau \in \mathbb{R}_+$ is a tunable hyperparameter. Typical candidates for these terms are $E(\mathbf{y}, \mathbf{H}(\mathbf{s})) = \|\mathbf{y} - \mathbf{H}(\mathbf{s})\|_2^2$ and $R(\mathbf{s}) = \|\mathbf{L}\mathbf{s}\|_p^p$ [1], [2], [3], [4], [5], [6], [7] with $p \in [1, 2]$. Here, \mathbf{L} is a linear transformation such as the discrete version of the wavelet transform or the gradient operator, which takes part in the regularization. For instance, total-variation (TV) regularization [3] uses the ℓ_1 -norm along with the gradient operator, which promotes solutions with sparse derivatives. It is widely used for compressed sensing and extreme imaging applications where the data is scarce [7]. The resulting optimization problems are typically solved by iterative algorithms such as gradient descent, the fast iterative shrinkage-thresholding algorithm (FISTA) [8], [9], [10], or the alternating-direction method of multipliers (ADMM) [11].

B. Bayesian Inference

In the Bayesian approach to image reconstruction [12], [13], [14], [15], the signal \mathbf{s} is modeled as the realization of a random vector with a suitable probability density function (pdf) p_S that captures our prior knowledge about the signal. The idea here is to characterize the posterior distribution

$$p_{S|Y}(\mathbf{s}|\mathbf{y}) \propto p_N(\mathbf{y} - \mathbf{H}(\mathbf{s})) p_S(\mathbf{s}), \quad (3)$$

which depends on the statistics of the noise p_N and on the prior distribution p_S , and to make inferences based on it.

The posterior distribution can be used for the derivation of several point estimators for the signal \mathbf{s} . One such example is the maximum a posteriori estimator, which is the mode of

the posterior distribution and leads to an optimization problem that resembles (2), with $E(\mathbf{y}, \mathbf{H}(\mathbf{s})) \propto (-\log(p_{\mathbf{N}}(\mathbf{y} - \mathbf{H}(\mathbf{s}))))$ and $R(\mathbf{s}) \propto (-\log(p_{\mathbf{S}}(\mathbf{s})))$, thus linking the variational and Bayesian approaches [16], [17], [18], [19]. Another example is the minimum mean-square error (MMSE) estimator which turns out to be the posterior mean [19].

Besides the derivation of point estimators, the Bayesian framework allows one to quantify the uncertainty of the reconstructed image. This feature offers an interesting perspective for computational imaging as most practical reconstruction schemes, including the variational ones, do not provide any assessment of reliability.

In general, inference tasks entail the estimation of expected values with respect to the posterior distribution. Typically, these are high-dimensional integrals that cannot be evaluated analytically. Thus, one relies on Markov chain Monte Carlo (MCMC) methods to efficiently draw samples from the posterior and then use them to approximate the integrals [20], [21], [22], [23].

C. Deep-Learning-Based Methods

Over the past few years, researchers have started to deploy deep-learning-based methods to solve inverse problems in imaging. The learning-based methods have been found to outperform the traditional model-based ones. Broadly speaking, their underlying principle is to utilize large amounts of training data to improve the reconstruction quality, as opposed to the specification of prior information about the image of interest in the form of mathematical models, as in the variational and Bayesian approaches described earlier.

The first generation of deep-learning-based methods involves training a convolutional neural network (CNN) as a nonlinear mapping that relates a low-quality estimate of the signal to the desired high-quality estimate [24], [25], [26], [27], [28]. The reconstruction pipeline then consists of using a fast classical algorithm to yield an initial solution and then correcting for its artifacts using the trained CNN. This category of methods includes “unrolling” [29], [30], [31], [32], [33], [34], where the architecture of the CNN is designed by studying iterations of algorithms used for solving Problem (2). While the first-generation end-to-end learning methods have achieved state-of-the-art performances in several inverse problems, recent works have highlighted their instability and lack of robustness [35], [36].

The second generation of deep-learning-based methods aims at the integration of CNNs into iterative reconstruction algorithms. The plug-and-play priors (PnP) [37] and regularization-by-denoising (RED) [38] frameworks are two successful examples that provide a way to carry out this integration. In PnP algorithms, the proximal operator that appears in the iterations of the proximal algorithms (FISTA, ADMM) is replaced by a generic denoiser which imposes an implicit prior on the signal. RED, by contrast, incorporates an explicit regularization term that is constructed with the help of the chosen denoiser. In the learning-based variants of these frameworks, one uses appropriately trained CNNs as the denoising routines [39], [40], [41], [42], [43], [44], [45]. Another example of such methods

is projected gradient descent where the projection operator is a trained neural network that projects onto the space of desired signals [46], [47], [48]. Unlike the first-generation methods, the second-generation ones enforce consistency between the reconstructed signal and the measurements. They are also more versatile as the CNN denoisers can be used for several inverse problems without the need for retraining. One obstacle to the deployment of these learning-based iterative schemes is that the Lipschitz constant of the CNNs must be controlled in order to ensure their convergence [49], [50], which is not straightforward and remains an active area of research [50], [51], [52].

One can also identify a third class of deep-learning-based methods that make use of deep generative models such as variational autoencoders (VAE) [53] and generative adversarial networks (GAN) [54]. These models include a generator network that maps a low-dimensional latent space to the high-dimensional signal space. They are trained to capture the statistics of the dataset and generate sample signals similar to those in the dataset. Once such a deep generative model has been successfully trained, its application to an inverse problem typically consists of finding the optimal latent variable such that the resulting signal best fits the measurements. Recent works have focused on the design and analysis of algorithms for the inversion of such generative models [55], [56], [57], [58], [59].

The three classes of deep-learning-based methods discussed so far are variational in nature and provide a single reconstruction as their output. The success of these methods has stimulated the development of Bayesian methods that exploit the power of neural networks. For instance, in [60], the authors propose two frameworks for “deep Bayesian inversion” that are analogues of the first generation end-to-end deep-learning-based methods and require training data consisting of signals and their corresponding measurements. Their first approach involves the training of a conditional GAN to sample from the posterior distribution, while their second approach deploys neural networks to approximate a chosen statistical estimator. More recently, the focus has been on the development of more modular Bayesian methods where only the prior is modeled by neural networks. This has led to various posterior sampling schemes for priors defined either implicitly through denoising CNNs (such as the ones used in the PnP or RED frameworks) [61], [62], [63] or through GANs [64], VAEs [65], [66], and score-based generative models [67], [68]. So far, most of these works have focused on inverse problems with linear or linearized forward models.

A current frontier in imaging is the inversion of nonlinear models, which arise in several applications, two notable examples being phase retrieval and optical diffraction tomography. Such applications could benefit greatly from the development of neural-network-based Bayesian reconstruction methods.

D. Contributions

In this paper, we present a Bayesian framework to solve a broad class of nonlinear inverse problems, where the prior is represented by a trained deep generative model. Our contributions are as follows.

- We develop a method based on the Metropolis-adjusted Langevin algorithm (MALA) [69], [70] to sample from the posterior distribution for the class of nonlinear inverse problems where the forward model has a neural-network-like structure. This class includes a wide variety of practical imaging modalities. We show that the structure of the forward model and the low-dimensional latent space of the generative prior enable tractable Bayesian inference.
- We introduce the concept of augmented generative models. This is motivated by the observation that deep generative models are easier to train when the dataset consists of images with the same range of pixel values. Unfortunately, such models are not well-matched to imaging modalities where one is interested in extracting the precise value of objects rather than merely visualizing contrast. Our proposed augmented models provide us with a simple but effective way of dealing with quantitative data.
- We illustrate the advantages of the proposed reconstruction framework through numerical experiments for two nonlinear imaging modalities: phase retrieval and optical diffraction tomography.

The paper is organized as follows: In Section II, we discuss the structure of the forward model for our nonlinear inverse problems. We detail the Bayesian reconstruction framework in Section III. There, we introduce augmented generative models and we explain our posterior-sampling scheme. We present our experimental results in Section IV.

II. NONLINEAR INVERSE PROBLEMS AND FORWARD MODELS

In this section, we start by describing the class of nonlinear inverse problems that we are interested in. We then focus on two concrete examples—phase retrieval and optical diffraction tomography—and detail the physical models involved.

A. Nonlinear Inverse Problems

The objective is to recover an image $\mathbf{s} \in \mathbb{R}^K$ from its noisy measurements $\mathbf{y} \in \mathbb{C}^M$ given by $\mathbf{y} = \mathbf{N}(\mathbf{y}_0)$ with

$$\mathbf{y}_0 = \mathbf{H}(\mathbf{s}), \quad (4)$$

where $\mathbf{H} : \mathbb{R}^K \rightarrow \mathbb{C}^M$ is a nonlinear operator that models the physics of the imaging system and $\mathbf{N} : \mathbb{C}^M \rightarrow \mathbb{C}^M$ is an operator that models the corruption of the measurements by noise. In this work, we consider the class of nonlinear forward models \mathbf{H} whose computational structure can be encoded by a directed acyclic graph and thus resembles a neural network.

The Jacobian matrix of \mathbf{H} at any point $\mathbf{x} = (x_1, \dots, x_K) \in \mathbb{R}^K$ is defined as

$$\mathbf{J}_{\mathbf{H}}(\mathbf{x}) = \begin{bmatrix} \frac{\partial}{\partial x_1} [\mathbf{H}(\mathbf{x})]_1 & \cdots & \frac{\partial}{\partial x_K} [\mathbf{H}(\mathbf{x})]_1 \\ \vdots & \ddots & \vdots \\ \frac{\partial}{\partial x_1} [\mathbf{H}(\mathbf{x})]_M & \cdots & \frac{\partial}{\partial x_K} [\mathbf{H}(\mathbf{x})]_M \end{bmatrix}. \quad (5)$$

Gradient-based MCMC methods (see Section III for a specific example) involve the computation of quantities such as $\mathbf{J}_{\mathbf{H}}^H(\mathbf{x})\mathbf{r}$ for some vectors $\mathbf{x} \in \mathbb{R}^K$, $\mathbf{r} \in \mathbb{C}^M$, and this can be a potential bottleneck. The neural-network-like structure of \mathbf{H} allows us to




Fig. 1. The forward model for phase retrieval (6) expressed as a one-layer fully-connected neural network with linear weights \mathbf{A} and quadratic activation functions.

compute these efficiently using the error backpropagation algorithm. This, in turn, makes Bayesian inference computationally feasible.

The class of nonlinear inverse problems that fit this description is very broad and adaptable to most existing imaging modalities. In principle, it covers all possible inverse problems, in particular, the linear case is trivially covered. More generally, if sufficient data is available, one can indeed train a neural network to mimic the physics of our forward model. Next, we look at two particular problems that nicely fall within our predefined class.

B. Phase Retrieval

Phase retrieval [71], [72] is a nonlinear inverse problem that is ubiquitous in computational imaging. It consists in the recovery of a signal from its intensity-only measurements and is a central issue in optics [73], [74], astronomy [75], [76], and computational microscopy [77], [78], [79], [80].

In the phase-retrieval problem that we consider in this paper, the noise-free measurements are modeled as

$$\mathbf{y}_0 = \mathbf{H}_{\text{pr}}(\mathbf{s}) = |\mathbf{A}\mathbf{s}|^2, \quad (6)$$

where $\mathbf{A} : \mathbb{R}^K \rightarrow \mathbb{C}^M$ is either the Fourier matrix [73], [80], [81] or some realization of a random matrix with independent and identically distributed (i.i.d.) elements [72], [82], [83], and where $|\cdot|^2$ is a component-wise operator. As shown in Fig. 1, the forward model in (6) can be expressed as a one-layer fully-connected neural network with fixed linear weights \mathbf{A} and quadratic activation functions.

C. Optical Diffraction Tomography

In optical diffraction tomography (ODT), the aim is to recover the refractive-index (RI) map of a sample from complex-valued measurements of the scattered fields generated when the sample is probed by a series of tilted incident fields [84]. According to the scalar-diffraction theory, the propagation of the incident fields through the sample is governed by the wave equation. While pioneering works relied on linear models to approximate this propagation [84], [85], recent works have significantly improved the quality of RI reconstruction by using more accurate nonlinear models that account for multiple scattering

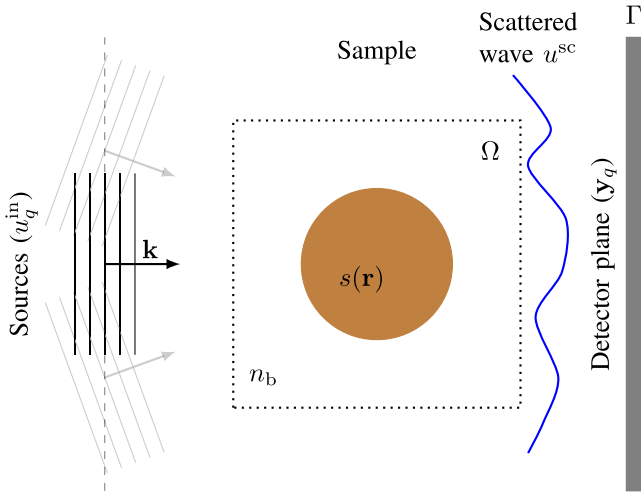


Fig. 2. Optical diffraction tomography. A sample of refractive index $n_b + s(\mathbf{r})$ is immersed in a medium of index n_b , and illuminated by an incident plane wave (wave vector \mathbf{k}). The interaction of the wave with the object produces scattered waves, which are recorded at the detector plane.

[86]. Here, we look at one such nonlinear model called the beam-propagation method (BPM).

Helmholtz Equation: We consider a sample with a real-valued spatially varying refractive index that is immersed in a medium with constant refractive index n_b , as shown in Fig. 2. The RI distribution in the region of interest $\Omega = [0, L_x] \times [0, L_z]$ is represented as $n(\mathbf{r}) = n_b + s(\mathbf{r})$, where $\mathbf{r} = (x, z)$ and $s(\mathbf{r})$ is the RI contrast. The sample is illuminated with an incident plane wave $u^{\text{in}}(\mathbf{r})$ of free-space wavelength λ , whose direction of propagation is specified by the wave vector \mathbf{k} . The total field $u(\mathbf{r})$ that results from the interaction between the sample and the incident wave is then recorded at the positions $\{\mathbf{r}_m\}_{m=1}^{M'}$ in the detector plane Γ to yield the complex measurements $\mathbf{y} \in \mathbb{C}^{M'}$. The interplay between the total field $u(\mathbf{r})$ at any point in space and the refractive index contrast $\delta n(\mathbf{r})$ is described by the Helmholtz equation

$$\nabla^2 u(\mathbf{r}) + k_0^2 n^2(\mathbf{r}) u(\mathbf{r}) = 0, \quad (7)$$

where $k_0 = \frac{2\pi}{\lambda}$.

Beam Propagation Method: For computational purposes, the region of interest Ω is subdivided into an $(N_x \times N_z)$ array of pixels with sampling steps δ_x and δ_z along the first and second dimension, respectively. The corresponding samples of the RI contrast $s(\mathbf{r})$ and total field $u(\mathbf{r})$ are stored in the vectors¹ $\mathbf{s} \in \mathbb{R}^K$ and $\mathbf{u} \in \mathbb{C}^K$, respectively, where $K = N_x N_z$. Further, let $\mathbf{s}_k \in \mathbb{R}^{N_x}$ and $\mathbf{u}_k \in \mathbb{C}^{N_x}$ represent the above quantities when restricted to the slice $z = k\delta_z$.

BPM computes the total field \mathbf{u} in a slice-by-slice manner along the z -axis. For a given incident wave $u^{\text{in}}(\mathbf{r})$ that is propagated over a region larger than Ω , we set the initial conditions as $\mathbf{u}_{-1}(\mathbf{s}) = (u^{\text{in}}(i\delta_x, -\delta_z))_{i=0}^{N_x-1} \in \mathbb{C}^{N_x}$. The total field over Ω is

¹Since the total field $u(\mathbf{r})$ depends on the RI contrast $s(\mathbf{r})$, we also refer to its discretized version as $\mathbf{u}(\mathbf{s})$.




Fig. 3. The computational structure for BPM resembles a neural network.

then computed via a series of diffraction and refraction steps

$$\tilde{\mathbf{u}}_k(\mathbf{s}) = \mathbf{u}_{k-1}(\mathbf{s}) * \mathbf{h}_{\text{prop}}^{\delta_z} \quad (\text{diffraction}) \quad (8)$$

$$\mathbf{u}_k(\mathbf{s}) = \tilde{\mathbf{u}}_k(\mathbf{s}) \odot \mathbf{p}_k(\mathbf{s}) \quad (\text{refraction}), \quad (9)$$

where $k = 0, 1, \dots, (N_z - 1)$, and the symbols $*$ and \odot stand for convolution and pointwise multiplication, respectively. The convolution kernel $\mathbf{h}_{\text{prop}}^{\delta_z} \in \mathbb{C}^{N_x}$ for the diffraction step is characterized in the Fourier domain as

$$\mathcal{F}\{\mathbf{h}_{\text{prop}}^{\delta_z}\}(\mathbf{w}_x) = e^{j\delta_z(\sqrt{k_0^2 n_b^2 - \mathbf{w}_x^2})}, \quad (10)$$

where \mathcal{F} denotes the discrete Fourier transform and $\mathbf{w}_x \in \mathbb{R}^{N_x}$ is the frequency variable. The subsequent refraction step involves a pointwise multiplication with the phase mask

$$\mathbf{p}_k(\mathbf{s}) = e^{jk_0 \delta_z \mathbf{s}_k}. \quad (11)$$

Finally, we define an operator $\mathbf{R} : \mathbb{C}^{N_x} \mapsto \mathbb{C}^{M'}$ that propagates $\mathbf{u}_{N_z-1}(\mathbf{s})$ to the detector plane Γ and restricts it to the sensor positions to give us the measurements $\mathbf{y} \in \mathbb{C}^{M'}$. Thus, for a given incident wave u^{in} , our noise-free nonlinear BPM forward model is of the form

$$\mathbf{y}_0 = \mathbf{H}_{\text{bpm}}(\mathbf{s}; u^{\text{in}}) = \mathbf{R}(\mathbf{u}_{N_z-1}(\mathbf{s})). \quad (12)$$

In Fig. 3, we show the implementation of \mathbf{H}_{bpm} as a directed acyclic graph.

Complete Forward Model: We assume that the sample is illuminated with Q incident plane waves $\{u_q^{\text{in}}\}_{q \in \{1, \dots, Q\}}$ and that the corresponding measurements are $\{\mathbf{y}_q \in \mathbb{C}^{M'}\}_{q \in \{1, \dots, Q\}}$. These measurements are related to the RI contrast \mathbf{s} of the sample through the BPM forward model in (12). We define a stacked measurement vector as $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_Q) \in \mathbb{R}^M$ ($M = QM'$). This allows us to rewrite the complete forward model in the form of (4), where the operator \mathbf{H} consists of the application of \mathbf{H}_{bpm} with all the illuminations and the concatenation of the outputs into a single vector.

III. BAYESIAN RECONSTRUCTION FRAMEWORK

We now present our reconstruction framework that is based on Bayesian statistics for solving the generic nonlinear inverse problem described in Section II-A. The image \mathbf{s} is assumed to be a realization of a random vector with pdf p_S and the statistical model for measurement noise is included within the likelihood function $p_{Y|S}$, which is the conditional distribution

of the measurements given the image. The quantity of interest here is the posterior distribution $p_{S|Y}$ as it provides a complete statistical characterization of the problem at hand. Using Bayes' rule, we then write $p_{S|Y}$ as

$$p_{S|Y}(\mathbf{s}|\mathbf{y}) = \frac{p_{Y|S}(\mathbf{y}|\mathbf{s})p_S(\mathbf{s})}{\int_{\mathbb{R}^K} p_{Y|S}(\mathbf{y}|\mathbf{s})p_S(\mathbf{s}) d\mathbf{s}}. \quad (13)$$

In this section, we first characterize the likelihood function $p_{Y|S}$. We then discuss the prior distribution p_S , which, in our framework, is defined through a deep generative model, followed by the posterior distribution $p_{S|Y}$. Finally, we detail a MCMC scheme to generate samples from the posterior distribution. This allows us to perform inference by computing point estimates and the uncertainties associated with them.

A. Likelihood Function

In our framework, we assume that the operator $\mathbf{N} : \mathbf{y}_0 \mapsto \mathbf{N}(\mathbf{y}_0)$ in (4) samples the noisy measurement vector \mathbf{y} from a conditional distribution $p_{Y|Y_0}$ according to

$$\mathbf{y} \sim p_{Y|Y_0}(\cdot|\mathbf{y}_0 = \mathbf{H}(\mathbf{s})), \quad (14)$$

where $p_{Y|Y_0}$ models the statistics of the noise in the imaging system. Since our forward models \mathbf{H} are deterministic, the quantity $p_{Y|S}$ (a.k.a. the likelihood function) is given by

$$p_{Y|S}(\mathbf{y}|\mathbf{s}) = p_{Y|Y_0}(\mathbf{y}|\mathbf{y}_0 = \mathbf{H}(\mathbf{s})). \quad (15)$$

In many imaging systems, there exist multiple independent sources of noise. It is therefore reasonable to assume an additive white-Gaussian-noise (AWGN) model, as dictated by the central limit theorem. There, the distribution $p_{Y|Y_0}$ is

$$p_{Y|Y_0}(\mathbf{y}|\mathbf{y}_0) \propto \exp\left(-\frac{\|\mathbf{y} - \mathbf{y}_0\|_2^2}{2\sigma^2}\right), \quad (16)$$

where σ is the standard deviation of the Gaussian noise.

Another model that is commonly used is the shot- or Poisson-noise model. In this case, we have that

$$p_{Y|Y_0}(\mathbf{y}|\mathbf{y}_0) = \prod_{m=1}^M \frac{([y_0]_m)^{[y]_m}}{([y]_m)!} \exp(-[y_0]_m), \quad (17)$$

where $\mathbf{y} \in \mathbb{N}^M$.

B. Prior Distribution

The choice of the distribution p_S reflects our prior knowledge about the image of interest. This information is crucial for the resolution of the inverse problem, especially when it is ill-posed. In classical Bayesian methods, p_S is chosen from a family of distributions with closed-form analytical expressions such that it fits the characteristics of the image and also allows for efficient inference. Popular examples include the Gaussian and Markovian models. In our framework, we instead propose to leverage the power of neural networks to define a data-driven prior distribution.

We assume that we have access to a dataset that contains sample images from the true (but unknown) probability distribution p_{image} of our image of interest. The idea then is to approximate p_{image} with p_S as defined by a deep generative model. More specifically, we consider generative models consisting of a generator network $G : \mathbb{R}^d \rightarrow \mathbb{R}^K$ ($d \ll K$) that maps a low-dimensional latent space to the high-dimensional image space. This network takes a vector $\mathbf{z} \in \mathbb{R}^d$, which is sampled from some distribution p_Z (typically a Gaussian or uniform distribution), and outputs a sample image $G(\mathbf{z})$. Thus, the generator network G and the distribution p_Z implicitly characterize p_S and provide us with a way to directly sample from it. If this model is properly trained, the resulting p_S is close to p_{image} and the generated images are statistically similar to the ones in the dataset.

In our experiments (see Section IV), we use the well-known Wasserstein GANs (WGANs) [87] for our data-driven prior. We provide a brief description of WGANs in the supplementary material.

Augmented Deep Generative Priors: The training of deep generative models such as GANs requires large amounts of data and is a challenging task in general. Over the past few years, there have been several proposals for performance improvements that have led to the development of better training schemes and network architectures. Most existing works use normalized datasets, where each image has the same range of pixel values. However, this is not suitable if we wish to use such models as priors in quantitative imaging (e.g., ODT). In these modalities, it is important to recover the actual values of the object (image) as compared to only the contrast. Thus, we require our generative model to be able to output images with different ranges of pixel values.

While performing our experiments, we observed that the training of high-quality WGANs on unnormalized datasets was non-trivial. We propose a simple effective workaround, which simplifies the training and allows us to build models that generate images with different ranges. We define an augmented generative model $G_h : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^K$ ($d \ll K$) that consists of a (standard) generative network $G : \mathbb{R}^d \rightarrow \mathbb{R}^K$ trained on a normalized dataset and a deterministic function $h : \mathbb{R} \rightarrow \mathbb{R}$. Here, the latent vector $\mathbf{z} = (\mathbf{z}_1, z_2) \in \mathbb{R}^{d+1}$ has two independent components $\mathbf{z}_1 \in \mathbb{R}^d$ and $z_2 \in \mathbb{R}$ that are sampled from p_{Z_1} and p_{Z_2} , respectively. The output image is given by $G_h(\mathbf{z}) = h(z_2)G(\mathbf{z}_1)$. For a generated image $G_h(\mathbf{z}) \in \mathbb{R}^K$, the term $G(\mathbf{z}_1) \in \mathbb{R}^K$ represents its details or contrast, and the term $h(z_2)$ represents its scaling factor. Since G is now required to only produce images with the same range, we can rely on existing GANs to obtain high-quality models. Moreover, the distribution of the scaling factor can be easily controlled by carefully choosing the distribution p_{Z_2} and the function h .

C. Posterior Distribution

Now that we are equipped with the likelihood function $p_{Y|S}$ and the prior distribution p_S , we look at the posterior distribution $p_{S|Y}$ of the image. Since our prior distribution p_S

is defined by a pre-trained augmented deep generative model $G_h: \mathbb{R}^{d+1} \rightarrow \mathbb{R}^K$, $\mathbf{z} \mapsto G_h(\mathbf{z})$ with $p_Z(\mathbf{z}) = p_{Z_1}(\mathbf{z}_1)p_{Z_2}(z_2)$ for any $\mathbf{z} = (\mathbf{z}_1, z_2) \in \mathbb{R}^{d+1}$, our $p_{S|Y}$ is given by the push-forward of the posterior distribution $p_{Z|Y}$ of the latent vector through the mapping G_h . The distribution $p_{Z|Y}$ can be written as

$$p_{Z|Y}(\mathbf{z}|\mathbf{y}) = \frac{p_{Y|Z}(\mathbf{y}|\mathbf{z})p_Z(\mathbf{z})}{\int_{\mathbb{R}^{d+1}} p_{Y|Z}(\mathbf{y}|\mathbf{z})p_Z(\mathbf{z}) d\mathbf{z}},$$

where $p_{Y|Z}(\mathbf{y}|\mathbf{z}) = p_{Y|S}(\mathbf{y}|\mathbf{s} = G_h(\mathbf{z}))$.

A Bayesian inverse problem is said to be well-posed in some metric on the space of probability measures if its solution (the posterior distribution) exists, is unique, and is continuous with respect to the measurements for the chosen metric [88]. Depending on the metric, the well-posedness of the Bayesian inverse problem ensures continuity of posterior expectations of appropriate quantities of interest. Based on the work in [88], we can show that for the AWGN model, our Bayesian problem is well-posed in the Prokhorov, total-variation and Hellinger distances. Moreover, our problem is well-posed in the Wasserstein distance if p_Z satisfies a finite-moment-like condition. By using a result from [66], we can also show the existence of the moments of our posterior distribution under mild conditions on p_Z and G_h . We provide the details regarding these properties in Appendix A.

D. Sampling From the Posterior Distribution

The proposed framework allows one to draw samples in the low-dimensional latent space instead of the high-dimensional image space directly. Specifically, if we generate a sample $\bar{\mathbf{z}}$ from $p_{Z|Y}$, then the image $\bar{\mathbf{s}} = G_h(\bar{\mathbf{z}})$ is a sample from $p_{S|Y}$.

In this work, we use the Metropolis-adjusted Langevin algorithm (MALA) [69], [70], which is a MCMC method, to sample from $p_{Z|Y}$. Given a sample $\bar{\mathbf{z}}_t$, MALA generates $\bar{\mathbf{z}}_{t+1}$ in two steps. In the first step, we construct a proposal $\tilde{\mathbf{z}}_{t+1}$ for the new sample according to

$$\tilde{\mathbf{z}}_{t+1} = \bar{\mathbf{z}}_t + \eta \nabla_{\mathbf{z}} \log p_{Z|Y}(\bar{\mathbf{z}}_t|\mathbf{y}) + \sqrt{2\eta}\zeta, \quad (18)$$

where ζ is drawn from the standard multivariate Gaussian distribution and $\eta \in \mathbb{R}_+$ is a fixed step-size. In the second step, the proposal $\tilde{\mathbf{z}}_{t+1}$ is either accepted or rejected, the acceptance probability being

$$\alpha = \min \left\{ 1, \frac{p_{Z|Y}(\tilde{\mathbf{z}}_{t+1}|\mathbf{y})q_Y(\bar{\mathbf{z}}_t|\tilde{\mathbf{z}}_{t+1})}{p_{Z|Y}(\bar{\mathbf{z}}_t|\mathbf{y})q_Y(\tilde{\mathbf{z}}_{t+1}|\bar{\mathbf{z}}_t)} \right\}, \quad (19)$$

where $q_Y(\bar{\mathbf{z}}|\tilde{\mathbf{z}}) = \exp(-\frac{1}{4\eta} \|\bar{\mathbf{z}} - \tilde{\mathbf{z}} - \eta \nabla_{\mathbf{z}} \log p_{Z|Y}(\tilde{\mathbf{z}}|\mathbf{y})\|_2^2)$. If the proposal is accepted, then we set $\bar{\mathbf{z}}_{t+1} = \tilde{\mathbf{z}}_{t+1}$; otherwise, $\bar{\mathbf{z}}_{t+1} = \bar{\mathbf{z}}_t$. One advantage of MALA is that it uses the gradient of the (log) target distribution to construct more probable proposals. In doing so, it explores the target distribution faster than some other MCMC methods such as the well-known random walk Metropolis-Hastings algorithm [89].

The major computational bottleneck in MALA is the computation of the gradient term $\nabla_{\mathbf{z}} \log p_{Z|Y}$ as it involves terms such as $\mathbf{J}_{\mathbf{H}}^H(\mathbf{x}_1)\mathbf{r}_1$ and $\mathbf{J}_{G_h}^H(\mathbf{x}_2)\mathbf{r}_2$, where $\mathbf{x}_1 \in \mathbb{R}^K$, $\mathbf{r}_1 \in \mathbb{C}^M$, $\mathbf{x}_2 \in \mathbb{R}^{d+1}$, and $\mathbf{r}_2 \in \mathbb{R}^K$. For instance, if we assume an AWGN

model with variance σ^2 and that p_Z is the standard multivariate Gaussian distribution, then $p_{Z|Y}$ can be written as

$$p_{Z|Y}(\mathbf{z}|\mathbf{y}) = \frac{1}{C} \exp \left(-\frac{\|\mathbf{y} - \mathbf{H}\{G_h(\mathbf{z})\}\|_2^2}{2\sigma^2} - \frac{\|\mathbf{z}\|_2^2}{2} \right), \quad (20)$$

where C is the normalization factor. In this case, the gradient term is

$$\begin{aligned} \nabla_{\mathbf{z}} \log p_{Z|Y}(\mathbf{z}|\mathbf{y}) &= -\frac{\mathbf{J}_{G_h}^H(\mathbf{z})\mathbf{J}_{\mathbf{H}}^H(G_h(\mathbf{z}))(\mathbf{y} - \mathbf{H}\{G_h(\mathbf{z})\})}{\sigma^2} - \mathbf{z}. \end{aligned} \quad (21)$$

Since G_h is a neural network and \mathbf{H} has a neural-network-like structure, we then compute $\nabla_{\mathbf{z}} \log p_{Z|Y}$ efficiently using an error backpropagation algorithm.

Once we have obtained the samples $\{\bar{\mathbf{z}}_t\}_{t=1}^T$ from $p_{Z|Y}$, we transform them to get the samples $\{G_h(\bar{\mathbf{z}}_t)\}_{t=1}^T$ from $p_{S|Y}$ and use them to perform inference. Specifically, we approximate any integral of the form $\int_{\mathbb{R}^K} f(\mathbf{s})p_{S|Y}(\mathbf{s}|\mathbf{y}) d\mathbf{s}$, where $f: \mathbb{R}^K \rightarrow \mathbb{R}$ is a real-valued function, by its empirical estimate $E_T(f) = \frac{1}{T} \sum_{t=1}^T f(G_h(\bar{\mathbf{z}}_t))$.

In practice, we discard some of the samples generated at the beginning of the chain to correct for their bias. This ‘‘burn-in’’ period can often be shortened by choosing a suitable starting point for the chain. We propose to initialize MALA with

$$\mathbf{z}_{\text{init}} = \arg \min_{\mathbf{z} \in \mathbb{R}^{d+1}} \|\mathbf{s}_{\text{init}} - G_h(\mathbf{z})\|_2^2, \quad (22)$$

where \mathbf{s}_{init} is a low-quality estimate obtained by using some fast classical reconstruction algorithm.

IV. RESULTS AND DISCUSSION

In this section, we show the benefits of our neural-network-based Bayesian reconstruction framework by applying it to both phase retrieval and optical diffraction tomography.

A. Augmented WGANs

In our first experiment, we highlight the importance of the proposed augmented generative models. We consider the task of training WGAN models on synthetic datasets consisting of (128×128) images, where each image contains a constant-valued disc and its background pixels are zero-valued. The coordinates (x, y) of the center of the disc, its radius r (in pixels), and its constant-intensity value v follow the uniform distributions $U_{(10,115)}$, $U_{(10,115)}$, $U_{[8,35]}$, and $U_{(0,0.2)}$, respectively. The aforementioned parameters implicitly define the probability distribution p_{data} that we wish to approximate using WGANs.

We qualitatively compare the performance of two models. The first model is a WGAN trained on 50,000 images sampled from p_{data} . In this case, the distribution p_Z for the latent variable is chosen to be the standard multivariate Gaussian distribution. The second model is an augmented WGAN, where the WGAN component is trained on a normalized dataset with 50,000 images. Thus, we first sample 50,000 images from p_{data} and we then normalize each of them such that the value of the disc is one. The distributions p_{Z_1} and p_{Z_2} are chosen to be standard




Fig. 4. Samples generated by trained models.

Gaussian distributions as well, and the function h is

$$h(x) = \frac{0.2}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{t^2}{2}} dt. \quad (23)$$

This choice of h and p_{Z_2} ensures that the scaling factor of the augmented WGAN follows the uniform distribution $U_{(0,0.2]}$. For both the models, we use the generator and critic network architectures described in the supplementary material. The WGAN is trained for 2500 epochs while the augmented WGAN is trained for 1250 epochs using RMSProp optimizers with a learning rate of 5×10^{-5} and a batch size of 64. The parameters λ_{gp} and n_{critic} (refer to the supplementary material) are set as 10 and 5, respectively.

In Fig. 4, we present typical samples generated by the two models. We observe that the augmented WGAN, unlike the WGAN, is able to produce sharp constant-valued discs.

B. Phase Retrieval

Next, we look at the phase-retrieval problem. We present two examples where the ground-truth images are taken from the MNIST [90] and Fashion-MNIST [91] testing datasets. In both cases, the measurements $\mathbf{y} \in \mathbb{N}^M$ are simulated according to (6) with a Poisson-noise model, where \mathbf{A} is one realization of a random matrix with i.i.d. entries from a zero-mean Gaussian distribution with variance $\sigma_{\mathbf{A}}^2$.

1) *MNIST*: The MNIST dataset contains (28×28) images of handwritten digits. The ground-truth image (Fig. 5) is first normalized to have values in the range $[0,1]$ and is then multiplied by a factor α which is picked uniformly at random from $(0, 0.5]$.

In this case, the WGAN component of our augmented model G_h is trained on the normalized MNIST training dataset which contains 50,000 images with values in the range $[0,1]$. The distributions p_{Z_1} and p_{Z_2} are standard Gaussian distributions

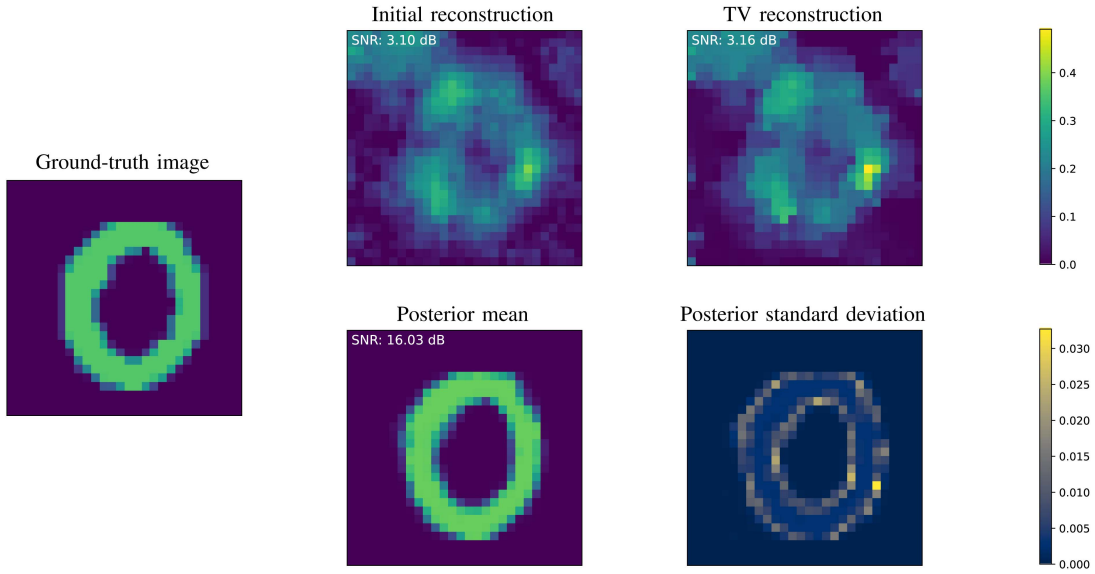


Fig. 5. Reconstructions for phase retrieval (oversampling ratio $M/K = 0.1$).

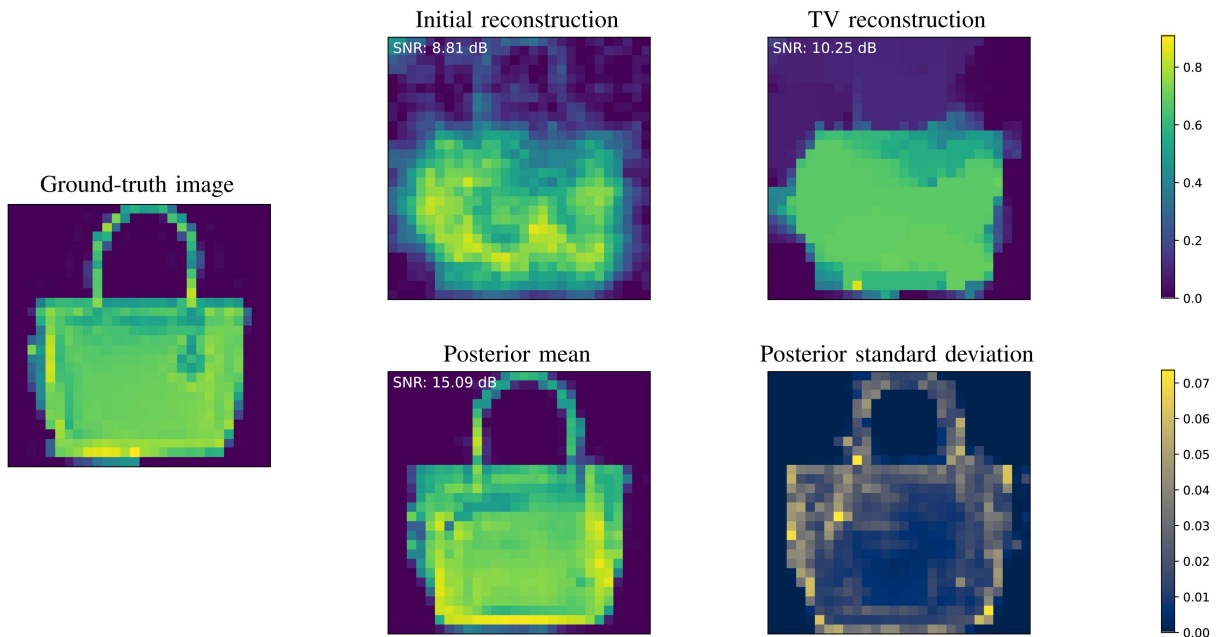


Fig. 6. Reconstructions for phase retrieval (oversampling ratio $M/K = 0.15$).

and the function h is

$$h(x) = \frac{0.5}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{t^2}{2}} dt. \quad (24)$$

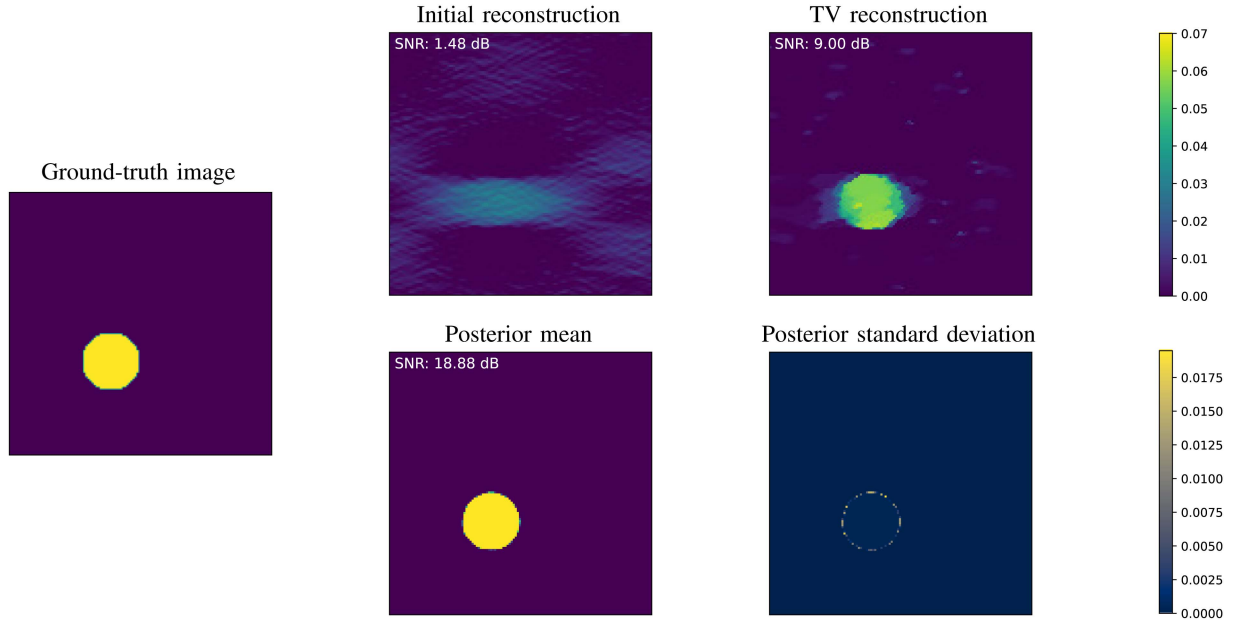
The architectures for the generator and critic networks can be found in the supplementary material. The WGAN is trained for 2000 epochs using ADAM optimizers [92] with a learning rate of 2×10^{-4} , hyperparameters $(\beta_1, \beta_2) = (0.5, 0.999)$, and a batch size of 64. The parameters λ_{gp} and n_{critic} are set as 10 and 5, respectively.

2) *Fashion-MNIST*: The Fashion-MNIST dataset consists of (28×28) grayscale images of different fashion products. Our ground-truth image from this dataset is shown in Fig. 6.

Here, the WGAN for our augmented deep generative prior is trained on the normalized Fashion-MNIST training dataset. It contains 60,000 images whose values lie in the range $[0, 1]$. The distributions p_{Z_1} and p_{Z_2} are taken as standard Gaussian distributions while the function h is

$$h(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{t^2}{2}} dt. \quad (25)$$

We provide the architectures for the generator and critic networks in the supplementary material. The WGAN is trained for 2250 epochs using ADAM optimizers with a learning rate of 2×10^{-4} , hyperparameters $(\beta_1, \beta_2) = (0.5, 0.999)$, and a batch

Fig. 7. Reconstructions for ODT ($v = 0.07$).

size of 64. The parameters λ_{gp} and n_{critic} are set as 10 and 5, respectively.

3) *Methods:* As discussed in Section III-D, we draw samples from the posterior distribution using MALA. The estimate \mathbf{s}_{init} that we use for initializing the chain is taken to be the solution of a variational problem with Tikhonov regularization, as in

$$\mathbf{s}_{init} = \arg \min_{\mathbf{s} \in \mathbb{R}^K} \left(\sum_{m=1}^M \left(-[\mathbf{y}]_m \log([\mathbf{A}\mathbf{s}]_m^2) + [\mathbf{A}\mathbf{s}]_m^2 \right) + \tau \|\nabla \mathbf{s}\|_{2,2}^2 + i_+(\mathbf{s}) \right). \quad (26)$$

There, $\nabla : \mathbb{R}^K \rightarrow \mathbb{R}^{K \times 2}$ is the gradient operator, $\|\cdot\|_{p,q}$ is the (ℓ_p, ℓ_q) -mixed norm defined as

$$\|\mathbf{x}\|_{p,q} \triangleq \left(\sum_{u=1}^U \left(\sum_{v=1}^V ([\mathbf{x}]_{u,v})^p \right)^{q/p} \right)^{1/q} \quad \forall \mathbf{x} \in \mathbb{R}^{U \times V}, \quad (27)$$

$\tau \in \mathbb{R}_+$ is the regularization parameter and the functional i_+ given by

$$i_+(\mathbf{s}) = \begin{cases} 0, & \mathbf{s} \in \mathbb{R}_+^K \\ +\infty, & \text{otherwise} \end{cases} \quad (28)$$

enforces the non-negativity constraint on the solution. The data-fidelity term in (26) corresponds to the negative log-likelihood under the Poisson-noise model. We solve the problem in (26) using a projected-gradient-descent algorithm. The regularization parameter τ so that it minimizes the mean-square error (MSE) with respect to the ground-truth is chosen via grid search.

After discarding the first T_b samples (burn-in period), we collect the next T samples for performing inference. We compute the posterior mean which corresponds to the minimum mean-square error (MMSE) estimate. Further, to quantify the

uncertainty associated with our estimation, we also compute the pixel-wise standard-deviation map.

We compare the performance of our GAN-based posterior-mean estimator with that of the TV-regularized method [3]

$$\mathbf{s}_{TV} = \arg \min_{\mathbf{s} \in \mathbb{R}^K} \left(\sum_{m=1}^M \left(-[\mathbf{y}]_m \log([\mathbf{A}\mathbf{s}]_m^2) + [\mathbf{A}\mathbf{s}]_m^2 \right) + \tau \|\nabla \mathbf{s}\|_{2,1} + i_+(\mathbf{s}) \right). \quad (29)$$

TV regularization is known to promote piecewise-constant solutions and is well-matched to our test images. We solve (29) using FISTA [10] initialized with \mathbf{s}_{init} . The regularization parameter τ is tuned for optimal MSE performance with the help of a grid search.

4) *Results:* To illustrate the advantage of our neural-network-based prior, we consider extreme imaging settings where the number of measurements M is very small. For the first case (Fig. 5), we have that $\alpha = 0.36$, $M/K = 0.1$, $\sigma_{\mathbf{A}}^2 = 10$, $\eta = 10^{-5}$, $T_b = 8 \times 10^5$, and $T = 12 \times 10^5$. The parameters for the second case (Fig. 6) are $M/K = 0.15$, $\sigma_{\mathbf{A}}^2 = 0.5$, $\eta = 1.75 \times 10^{-6}$, $T_b = 17.5 \times 10^5$, and $T = 5 \times 10^5$.

In Figures 5 and 6, we see that the GAN-based posterior-mean estimator outperforms the TV-regularized method considerably. Here, the very low oversampling ratios severely affect the performance of TV regularization, even though it is a good fit for the underlying images. By contrast, despite the scarcity of measurements, our estimator remarkably yields excellent results. This highlights the potential of learning-based priors for highly ill-posed problems. Finally, we observe that, as one would expect, the standard-deviation maps indicate higher uncertainty at the edges for the posterior-mean estimator.

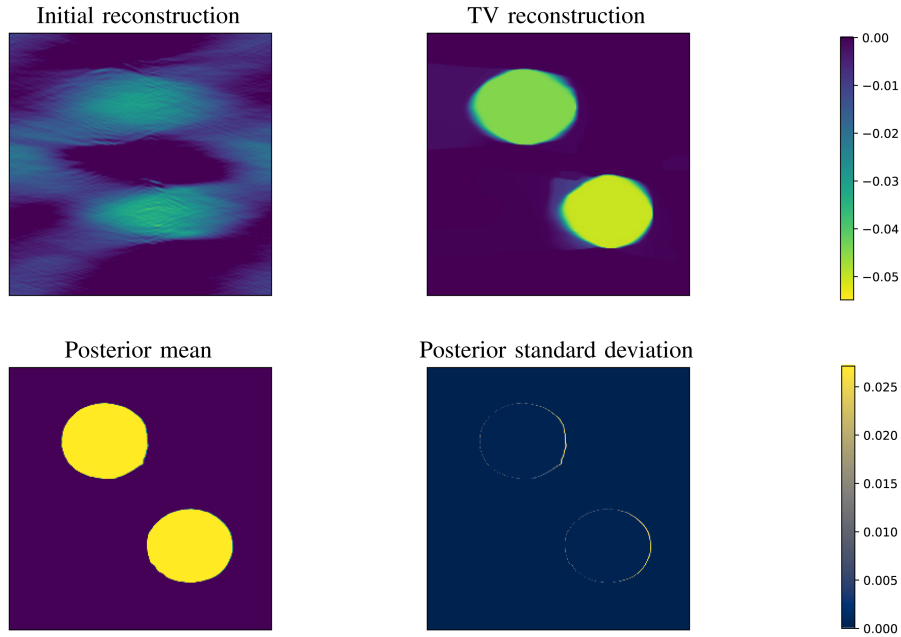


Fig. 8. Reconstructions for ODT (real data).

C. Optical Diffraction Tomography

We consider both simulated and real data for our ODT experiments.

1) *Simulated data*: In our simulated setup, the test image (Fig. 7) that represents the RI contrast is a random sample from the dataset described in Section IV-A: a disc with constant intensity v .

The measurements are simulated using the BPM of Section II-C with an AWGN model of variance $\sigma_n^2 = 0.05$. We set the sampling steps to $\delta_x = \delta_y = 0.1\mu\text{m}$, the medium RI to $n_b = 1.52$, and the wavelength to $\lambda = 0.406\mu\text{m}$. We use $Q = 20$ incident tilted plane waves with angles that are uniformly spaced in the range $[-\pi/12, \pi/12]$.

For this setting, we use the augmented WGAN prior of Section IV-A in our reconstruction framework.

2) *Real Data*: In our experiment with real data, the sample is a 2D cross-section of two non-overlapping fibres immersed in oil ($n_b = 1.525$) [93]. The RI contrast of the sample is negative. A standard Mach-Zehnder interferometer relying on off-axis digital holography ($\lambda = 0.450\mu\text{m}$) is used to collect measurements from $Q = 59$ views in the range $[-\pi/6, \pi/6]$.

We crop the acquired data such that the measurement vector for each view is of length $M' = 256$. We take the discretized region of interest to be of the size (256×256) and we set the sampling steps for BPM (used for reconstruction) to $\delta_x = \delta_y = 0.1257\mu\text{m}$. We assume an AWGN model of variance $\sigma_n^2 = 0.15$ for the measurements.

Here, the WGAN for our prior is trained on a synthetic dataset containing 100,000 images of size (256×256) , where each image consists of two non-overlapping discs with a constant intensity of one and a zero-valued background. The coordinates of the centers of the two discs are sampled from $U_{(20,235)}$ and their radii are sampled from $U_{[10,50]}$ subject to the constraint that

they do not overlap. The distributions p_{Z_1} and p_{Z_2} are standard Gaussian distributions and the function h is taken to be

$$h(x) = -\frac{0.1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{t^2}{2}} dt. \quad (30)$$

The architectures for the generator and critic networks are detailed in the supplementary material. The WGAN is trained for 500 epochs using RMSProp optimizers with a learning rate of 5×10^{-5} and a batch size of 128. The parameters λ_{gp} and n_{critic} are set as 10 and 5, respectively.

3) *Methods*: For both settings, the estimate \mathbf{s}_{init} for MALA is obtained by the application of a filtered backpropagation algorithm that uses the Rytov approximation [85] to model the scattering. We collect T samples from the posterior distribution using MALA with a step-size τ and burn-in period T_b , and use them to compute the posterior mean and pixel-wise standard-deviation map.

We compare our estimator with the TV-based method

$$\mathbf{s}_{\text{TV}} = \arg \min_{\mathbf{s} \in \mathbb{R}^K} \left(\sum_{q=1}^Q \|\mathbf{y}_q - \mathbf{H}_{\text{bpm}}(\mathbf{s}; u_q^{\text{in}})\|_2^2 + \tau \|\nabla \mathbf{s}\|_{2,1} + \mathcal{I}(\mathbf{s}) \right), \quad (31)$$

where $\mathcal{I}(\mathbf{s}) = i_+(\mathbf{s})$ for the simulated data and $\mathcal{I}(\mathbf{s}) = i_-(\mathbf{s})$ for the real data. This is a state-of-the-art method for ODT and is commonly used in practice [7], [94]. Moreover, it is well-suited for the constant-valued discs in our samples. The problem in (31) is solved using FISTA initialized with \mathbf{s}_{init} . The regularization parameter τ is tuned for optimal MSE performance in the simulated-data setting via a grid search, while it is tuned manually in the real-data setting.

4) *Results*: The settings that we consider for our ODT experiments are highly ill-posed as the incident waves only explore

a limited range. As a result, the measurements lack information along the horizontal axis, which leads to the so-called missing-cone problem. For the first case (Fig. 7), we have that $v = 0.07$, $\eta = 2 \times 10^{-7}$, $T_b = 2 \times 10^4$, and $T = 8 \times 10^4$. For the second case (Fig. 8), we have that $\eta = 5 \times 10^{-8}$, $T_b = 15 \times 10^4$, and $T = 5 \times 10^4$.

In Figs. 7 and 8, we observe that the TV reconstructions (and the initial ones) are elongated in the horizontal direction due to the lack of information along this axis. However, the GAN-based estimator is able to overcome the missing-cone problem. It yields reconstructions whose quality is remarkable.

D. Discussion

With the help of the above-described experiments, we have demonstrated the potential of our deep-generative-prior-based Bayesian reconstruction framework for challenging nonlinear inverse problems. We now mention some directions for future work which can further improve this framework.

In the present form, our scheme lacks theoretical guarantees for MALA to be geometrically ergodic (convergence to the equilibrium distribution at a geometric rate). A topic of future work could be to investigate the imposition of appropriate constraints on the generative model such that the resulting posterior distribution satisfies certain smoothness and tail conditions [95] that ensure geometric ergodicity of MALA.

The performance of our scheme heavily relies on how well the prior models the object of interest. Thus, any progress on the side of designing and training high-quality large-scale deep generative models could be translated to our framework.

While the neural-network-like structure of our forward models make our approach tractable, like MCMC methods in general, it requires a lot of computation. It could be interesting to consider alternatives to MALA that might help in speeding up this approach.

V. CONCLUSION

We have presented a Bayesian reconstruction framework for nonlinear inverse problems where the prior information on the image of interest is encoded by a deep generative model. Specifically, we have designed a tractable posterior-sampling scheme based on the Metropolis-adjusted Langevin algorithm for the class of nonlinear inverse problems where the forward model has a neural-network-like computational structure. This class includes most practical imaging modalities. We have proposed the concept of augmented generative models. They allow us to tackle the problem of the quantitative recovery of images. Finally, we have illustrated the benefits of our framework by applying it to two nonlinear imaging modalities—phase retrieval and optical diffraction tomography.

APPENDIX A

PROPERTIES OF THE POSTERIOR DISTRIBUTION

A. Well-Posedness

A Bayesian inverse problem is said to be well-posed in some metric on the space of probability measures if the posterior

distribution exists, is unique, and is continuous with respect to the measurements for the chosen metric [88]. Here, we present sufficient conditions from [88, Assumptions 3.5, 3.10 and Theorems 3.6, 3.12] that guarantee the well-posedness of our problem in the latent space, that is, with respect to $p_{Z|Y}$ as described in Section III-C.

The following conditions are stated for p_Z -almost every (a.e.) $\mathbf{z} \in \mathbb{R}^{d+1}$ and every $\mathbf{y} \in \mathbb{R}^M$.

Conditions:

- 1) $p_{Y|Z}(\cdot|\mathbf{z})$ is a strictly positive pdf.
- 2) $\int_{\mathbb{R}^{d+1}} |p_{Y|Z}(\mathbf{y}|\mathbf{z}')| p_Z(\mathbf{z}') d\mathbf{z}' < \infty$
- 3) There exists g with $\int_{\mathbb{R}^{d+1}} |g(\mathbf{z}')| p_Z(\mathbf{z}') d\mathbf{z}' < \infty$ such that $p_{Y|Z}(\mathbf{y}'|\cdot) \leq g$ for all $\mathbf{y}' \in \mathbb{R}^M$.
- 4) $p_{Y|Z}(\cdot|\mathbf{z})$ is continuous.
- 5) There exists g' with $\int_{\mathbb{R}^{d+1}} |g'(\mathbf{z}')| p_Z(\mathbf{z}') d\mathbf{z}' < \infty$ such that $\|\mathbf{z}''\|_2^p p_{Y|Z}(\mathbf{y}'|\mathbf{z}'') \leq g'(\mathbf{z}'')$, where $p \in [1, \infty)$, for p_Z -a.e. $\mathbf{z}'' \in \mathbb{R}^{d+1}$ and all $\mathbf{y}' \in \mathbb{R}^M$.

If the conditions (1) – (4) hold, our Bayesian inverse problem in the latent space is well-posed in the Prokhorov, Hellinger and total-variation distances. In addition, if condition (5) holds, then the problem is also well-posed in the Wasserstein p -distance.

For additive white-Gaussian-noise (AWGN) models, the conditions (1) – (4) are satisfied for any physical forward model \mathbf{H} and prior distribution p_Z . Further, if p_Z is such that $\int \|\mathbf{z}'\|_2^p p_Z(\mathbf{z}') d\mathbf{z}' < \infty$ (e.g., Gaussian distribution), condition (5) is also satisfied [88, Corollary 5.1]. As for the Poisson-noise models used in some of our experiments, they do not fall within this framework of well-posedness developed in [88].

B. Existence of Moments

Based on Proposition 3.6 in [66], we also present some conditions under which the moments of our posterior distribution $p_{S|Y}$ exist. If the augmented deep generative prior G_h is Lipschitz-continuous and the prior distribution p_Z has finite moments $\mathbb{E}_{p_Z}[\|\mathbf{z}\|^k]$ for $k = 1, 2, \dots, K$, then the K th posterior moment $\mathbb{E}_{p_{S|Y}}[\|\mathbf{s}\|^K]$ exists for almost all measurements \mathbf{y} .

The typical choice for p_Z is the standard Gaussian distribution, which has finite moments. The Lipschitz-continuity of G_h is guaranteed if the generative network G and the function h are both Lipschitz-continuous and bounded. The Lipschitz condition on the network G holds when its weights and biases are finite-valued and it consists of Lipschitz-continuous activation functions (e.g., ReLU, sigmoid). The boundedness of G is ensured when the activation function in the output layer is bounded (such as the sigmoid function). These are conditions that are satisfied by the networks used in Section IV. Further, in our experiments, we choose the function h to be a scaled version of the cumulative density function of the standard normal distribution, which is Lipschitz-continuous and bounded.

ACKNOWLEDGMENT

We would like to thank Dr. Joowon Lim and Prof. Demetri Psaltis for providing us with real data for optical diffraction tomography, and Dr. Aleix Boquet-Pujadas for helpful discussions.

REFERENCES

- [1] A. N. Tikhonov, "Solution of incorrectly formulated problems and the regularization method," *Sov. Math.*, vol. 4, pp. 1035–1038, 1963.
- [2] M. Bertero and P. Boccacci, *Introduction to Inverse Problems in Imaging*. Boca Raton, FL, USA: CRC Press, 1998.
- [3] L. I. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1–4, pp. 259–268, 1992.
- [4] M. Lustig, D. L. Donoho, and J. M. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magn. Reson. Med.: An Official J. Int. Soc. Magn. Reson. Med.*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [5] M. Figueiredo, R. Nowak, and S. Wright, "Gradient projection for sparse reconstruction: Application to compressed sensing and other inverse problems," *IEEE J. Sel. Topics Signal Process.*, vol. 1, no. 4, pp. 586–597, Dec. 2007.
- [6] J. Lim et al., "Comparative study of iterative reconstruction algorithms for missing cone problems in optical diffraction tomography," *Opt. Exp.*, vol. 23, no. 13, pp. 16933–16948, 2015.
- [7] U. S. Kamilov et al., "Optical tomographic image reconstruction based on beam propagation and sparse regularization," *IEEE Trans. Comput. Imag.*, vol. 2, no. 1, pp. 59–70, Mar. 2016.
- [8] M. A. Figueiredo and R. D. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Trans. Image Process.*, vol. 12, no. 8, pp. 906–916, Aug. 2003.
- [9] I. Daubechies, M. Defrise, and C. De Mol, "An iterative thresholding algorithm for linear inverse problems with a sparsity constraint," *Commun. Pure Appl. Math.: A. J. Issued Courant Inst. Math. Sci.*, vol. 57, no. 11, pp. 1413–1457, 2004.
- [10] A. Beck and M. Teboulle, "A fast iterative shrinkage-thresholding algorithm for linear inverse problems," *SIAM J. Imag. Sci.*, vol. 2, no. 1, pp. 183–202, 2009.
- [11] S. Boyd, N. Parikh, and E. Chu, *Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers*. Delft, The Netherlands: Now Publishers Inc, 2011.
- [12] A. Mohammad-Djafari, "A full Bayesian approach for inverse problems," in *Maximum Entropy and Bayesian Methods*. Dordrecht: Springer, 1996, pp. 135–144.
- [13] A. Mohammad-Djafari, "Bayesian inference for inverse problems," in *Proc. AIP Conf. Proc.*, 2002, vol. 617, pp. 477–496.
- [14] A. M. Stuart, "Inverse problems: A Bayesian perspective," *Acta Numerica*, vol. 19, pp. 451–559, 2010.
- [15] M. Dashti and A. M. Stuart, "The Bayesian approach to inverse problems," in *Handbook of Uncertainty Quantification*, Springer, 2017, pp. 311–428.
- [16] S. D. Babacan, R. Molina, and A. K. Katsagelos, "Bayesian compressive sensing using Laplace priors," *IEEE Trans. Image Process.*, vol. 19, no. 1, pp. 53–63, Jan. 2010.
- [17] M. Unser and P. D. Tafti, "Stochastic models for sparse and piecewise-smooth signals," *IEEE Trans. Signal Process.*, vol. 59, no. 3, pp. 989–1006, Mar. 2011.
- [18] R. Gribonval, "Should penalized least squares regression be interpreted as maximum a posteriori estimation?," *IEEE Trans. Signal Process.*, vol. 59, no. 5, pp. 2405–2410, May 2011.
- [19] M. Pereyra, "Revisiting maximum-a-posteriori estimation in log-concave models," *SIAM J. Imag. Sci.*, vol. 12, no. 1, pp. 650–670, 2019.
- [20] W. R. Gilks, S. Richardson, and D. Spiegelhalter, *Markov Chain Monte Carlo in Practice*. Boca Raton, FL, USA: Crc Press, 1995.
- [21] C. J. Geyer, "Practical Markov chain Monte Carlo," *Stat. Sci.*, pp. 473–483, 1992.
- [22] M. Pereyra et al., "A survey of stochastic simulation and optimization methods in signal processing," *IEEE J. Sel. Topics Signal Process.*, vol. 10, no. 2, pp. 224–241, Mar. 2016.
- [23] J. Kaipio and E. Somersalo, *Statistical and Computational Inverse Problems*. Berlin, Germany: Springer, 2006, vol. 160.
- [24] K. H. Jin, M. T. McCann, E. Froustey, and M. Unser, "Deep convolutional neural network for inverse problems in imaging," *IEEE Trans. Image Process.*, vol. 26, no. 9, pp. 4509–4522, Sep. 2017.
- [25] H. Chen et al., "Low-dose CT with a residual encoder-decoder convolutional neural network," *IEEE Trans. Med. Imag.*, vol. 36, no. 12, pp. 2524–2535, Dec. 2017.
- [26] C. M. Hyun, H. P. Kim, S. M. Lee, S. Lee, and J. K. Seo, "Deep learning for undersampled MRI reconstruction," *Phys. Med. Biol.*, vol. 63, no. 13, 2018, Art. no. 135007.
- [27] K. Monakhova, J. Yurtsever, G. Kuo, N. Antipa, K. Yanny, and L. Waller, "Learned reconstructions for practical mask-based lensless imaging," *Opt. Exp.*, vol. 27, no. 20, pp. 28075–28090, 2019.
- [28] D. Perdios, M. Vonlanthen, F. Martinez, M. Arditi, and J.-P. Thiran, "CNN-based image reconstruction method for ultrafast ultrasound imaging," *IEEE Trans. Ultrason., Ferroelect., Freq. Control*, vol. 69, no. 4, pp. 1154–1168, Apr. 2022.
- [29] K. Gregor and Y. LeCun, "Learning fast approximations of sparse coding," in *Proc. 27th Int. Conf. Mach. Learn.*, 2010, pp. 399–406.
- [30] Y. Chen and T. Pock, "Trainable nonlinear reaction diffusion: A flexible framework for fast and effective image restoration," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1256–1272, Jun. 2017.
- [31] Y. Yang, J. Sun, H. Li, and Z. Xu, "Deep ADMM-Net for compressive sensing MRI," in *Proc. 30th Int. Neural Inf. Process. Syst.*, 2016, pp. 10–18.
- [32] H. K. Aggarwal, M. P. Mani, and M. Jacob, "MoDL: Model-based deep learning architecture for inverse problems," *IEEE Trans. Med. Imag.*, vol. 38, no. 2, pp. 394–405, Feb. 2019.
- [33] J. Adler and O. Öktem, "Learned primal-dual reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1322–1332, Jun. 2018.
- [34] V. Monga, Y. Li, and Y. C. Eldar, "Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing," *IEEE Signal Process. Mag.*, vol. 38, no. 2, pp. 18–44, Mar. 2021.
- [35] V. Antun, F. Renna, C. Poon, B. Adcock, and A. C. Hansen, "On instabilities of deep learning in image reconstruction and the potential costs of AI," *Proc. Nat. Acad. Sci.*, vol. 117, no. 48, pp. 30088–30095, 2020.
- [36] N. M. Gottschling, V. Antun, B. Adcock, and A. C. Hansen, "The troublesome kernel: Why deep learning for inverse problems is typically unstable," 2020, *arXiv:2001.01258*.
- [37] S. V. Venkatakrisnan, C. A. Bouman, and B. Wohlberg, "Plug-and-play priors for model based reconstruction," in *Proc. IEEE Glob. Conf. Signal Inf. Process.*, 2013, pp. 945–948.
- [38] Y. Romano, M. Elad, and P. Milanfar, "The little engine that could: Regularization by denoising (RED)," *SIAM J. Imag. Sci.*, vol. 10, no. 4, pp. 1804–1844, 2017.
- [39] T. Tirer and R. Giryes, "Image restoration by iterative denoising and backward projections," *IEEE Trans. Image Process.*, vol. 28, no. 3, pp. 1220–1234, Mar. 2019.
- [40] E. Ryu, J. Liu, S. Wang, X. Chen, Z. Wang, and W. Yin, "Plug-and-play methods provably converge with properly trained denoisers," in *Proc. Int. Conf. Mach. Learn.*, 2019, pp. 5546–5557.
- [41] Y. Sun, J. Liu, and U. Kamilov, "Block coordinate regularization by denoising," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, pp. 382–392.
- [42] Z. Wu, Y. Sun, A. Matlock, J. Liu, L. Tian, and U. S. Kamilov, "SIMBA: Scalable inversion in optical tomography using deep denoising priors," *IEEE J. Sel. Topics Signal Process.*, vol. 14, no. 6, pp. 1163–1175, Oct. 2020.
- [43] J. Liu, Y. Sun, C. Eldeniz, W. Gan, H. An, and U. S. Kamilov, "RARE: Image reconstruction using deep priors learned without groundtruth," *IEEE J. Sel. Topics Signal Process.*, vol. 14, no. 6, pp. 1088–1099, Oct. 2020.
- [44] K. Zhang, Y. Li, W. Zuo, L. Zhang, L. Van Gool, and R. Timofte, "Plug-and-play image restoration with deep denoiser prior," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 44, no. 10, pp. 6360–6376, Oct. 2022.
- [45] Y. Sun, Z. Wu, X. Xu, B. Wohlberg, and U. S. Kamilov, "Scalable plug-and-play ADMM with convergence guarantees," *IEEE Trans. Comput. Imag.*, vol. 7, pp. 849–863, 2021.
- [46] J. R. Chang, C.-L. Li, B. Poczos, B. V. Kumar, and A. C. Sankaranarayanan, "One network to solve them all—solving linear inverse problems using deep projection models," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 5888–5897.
- [47] H. Gupta, K. H. Jin, H. Q. Nguyen, M. T. McCann, and M. Unser, "CNN-based projected gradient descent for consistent CT image reconstruction," *IEEE Trans. Med. Imag.*, vol. 37, no. 6, pp. 1440–1453, Jun. 2018.
- [48] F. Yang, T.-A. Pham, H. Gupta, M. Unser, and J. Ma, "Deep-learning projector for optical diffraction tomography," *Opt. Exp.*, vol. 28, no. 3, pp. 3905–3921, 2020.
- [49] H. H. Bauschke and P. L. Combettes, *Convex Analysis and Monotone Operator Theory in Hilbert Spaces* (CMS Books in Mathematics Series), 2nd ed. Cham, Switzerland: Springer, 2017.
- [50] J. Hertrich, S. Neumayer, and G. Steidl, "Convolutional proximal neural networks and plug-and-play algorithms," *Linear Algebra Appl.*, vol. 631, pp. 203–234, 2021.
- [51] M. Terris, A. Repetti, J.-C. Pesquet, and Y. Wiaux, "Building firmly nonexpansive convolutional neural networks," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2020, pp. 8658–8662.

- [52] P. Bohra, D. Perdios, A. Goujon, S. Emery, and M. Unser, "Learning Lipschitz-controlled activation functions in neural networks for plug-and-play image reconstruction methods," in *Proc. NeurIPS Workshop Deep Learn. Inverse Problems*, 2021. [Online]. Available: <https://openreview.net/forum?id=efCsbTzQTbH>
- [53] D. P. Kingma and M. Welling, "Auto-encoding variational Bayes," in *Proc. Int. Conf. Learn. Representations*, 2014.
- [54] I. Goodfellow et al., "Generative adversarial nets," in *Proc. 27th Int. Neural Inform. Process. Syst.*, 2014, pp. 2672–2680.
- [55] A. Bora, A. Jalal, E. Price, and A. G. Dimakis, "Compressed sensing using generative models," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 537–546.
- [56] V. Shah and C. Hegde, "Solving linear inverse problems using GAN priors: An algorithm with provable guarantees," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2018, pp. 4609–4613.
- [57] P. Hand, O. Leong, and V. Voroninski, "Phase retrieval under a generative prior," in *Proc. 32nd Int. Neural Inf. Process. Syst.*, 2018, pp. 9154–9164.
- [58] W. Huang, P. Hand, R. Heckel, and V. Voroninski, "A provably convergent scheme for compressive sensing under random generative priors," *J. Fourier Anal. Appl.*, vol. 27, no. 2, pp. 1–34, 2021.
- [59] M. González, A. Almansa, and P. Tan, "Solving inverse problems by joint posterior maximization with autoencoding prior," *SIAM J. Imag. Sci.*, vol. 15, no. 2, pp. 822–859, 2022.
- [60] J. Adler and O. Öktem, "Deep Bayesian inversion," 2018, *arXiv:1811.05910*.
- [61] Z. Kadkhodaie and E. P. Simoncelli, "Solving linear inverse problems using the prior implicit in a denoiser," in *Proc. NeurIPS Workshop Deep Learn. Inverse Problems*, 2020. [Online]. Available: <https://openreview.net/forum?id=RLN7K4U3UST>
- [62] B. Kawar, G. Vaksman, and M. Elad, "Snips: Solving noisy inverse problems stochastically," in *Proc. Adv. Neural Inform. Process. Syst.*, 2021, vol. 34, pp. 21757–21769.
- [63] R. Laumont, V. D. Bortoli, A. Almansa, J. Delon, A. Durmus, and M. Pereyra, "Bayesian imaging using plug & play priors: When langevin meets tweedie," *SIAM J. Imag. Sci.*, vol. 15, no. 2, pp. 701–737, 2022.
- [64] D. Patel and A. A. Oberai, "Bayesian inference with generative adversarial network priors," 2019, *arXiv:1907.09987*.
- [65] K. C. Tezcan, N. Karani, C. F. Baumgartner, and E. Konukoglu, "Sampling possible reconstructions of undersampled acquisitions in MR imaging with a deep learned prior," *IEEE Trans. Med. Imag.*, vol. 41, no. 7, pp. 1885–1896, Jul. 2022.
- [66] M. Holden, M. Pereyra, and K. C. Zygalakis, "Bayesian imaging with data-driven priors encoded by neural networks," *SIAM J. Imag. Sci.*, vol. 15, no. 2, pp. 892–924, 2022.
- [67] A. Jalal, M. Arvinte, G. Daras, E. Price, A. G. Dimakis, and J. Tamir, "Robust compressed sensing MRI with deep generative priors," *Adv. Neural Inf. Process. Syst.*, vol. 34, pp. 14938–14954, 2021.
- [68] Y. Song, L. Shen, L. Xing, and S. Ermon, "Solving inverse problems in medical imaging with score-based generative models," in *Proc. Int. Conf. Learn. Representations*, 2022. [Online]. Available: <https://openreview.net/forum?id=vaRCHVj0uGI>
- [69] G. O. Roberts and R. L. Tweedie, "Exponential convergence of Langevin distributions and their discrete approximations," *Bernoulli*, pp. 341–363, 1996.
- [70] G. O. Roberts and O. Stramer, "Langevin diffusions and Metropolis-Hastings algorithms," *Methodol. Comput. Appl. Probability*, vol. 4, no. 4, pp. 337–357, 2002.
- [71] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev, "Phase retrieval with application to optical imaging: A contemporary overview," *IEEE Signal Process. Mag.*, vol. 32, no. 3, pp. 87–109, May 2015.
- [72] F. Fogel, I. Waldspurger, and A. d'Aspremont, "Phase retrieval for imaging problems," *Math. Program. Comput.*, vol. 8, no. 3, pp. 311–335, 2016.
- [73] R. P. Millane, "Phase retrieval in crystallography and optics," *JOSA A*, vol. 7, no. 3, pp. 394–411, 1990.
- [74] A. M. Maiden and J. M. Rodenburg, "An improved ptychographical phase retrieval algorithm for diffractive imaging," *Ultramicroscopy*, vol. 109, no. 10, pp. 1256–1262, 2009.
- [75] J. R. Fienup, J. C. Marron, T. J. Schulz, and J. H. Seldin, "Hubble space telescope characterized by using phase-retrieval algorithms," *Appl. Opt.*, vol. 32, no. 10, pp. 1747–1767, 1993.
- [76] W. L. Freedman et al., "Final results from the Hubble space telescope key project to measure the Hubble constant," *Astrophysical J.*, vol. 553, no. 1, pp. 47–72, 2001.
- [77] M. H. Maleki and A. J. Devaney, "Phase-retrieval and intensity-only reconstruction algorithms for optical diffraction tomography," *JOSA A*, vol. 10, no. 5, pp. 1086–1092, 1993.
- [78] T. E. Gureyev and K. A. Nugent, "Rapid quantitative phase imaging using the transport of intensity equation," *Opt. Commun.*, vol. 133, no. 1–6, pp. 339–346, 1997.
- [79] F. Zernike, "Phase contrast, A new method for the microscopic observation of transparent objects—Part II," *Physica*, vol. 9, no. 10, pp. 974–986, 1942.
- [80] G. Zheng, R. Horstmeyer, and C. Yang, "Wide-field, high-resolution Fourier ptychographic microscopy," *Nature Photon.*, vol. 7, no. 9, pp. 739–745, 2013.
- [81] J. M. Rodenburg and H. M. Faulkner, "A phase retrieval algorithm for shifting illumination," *Appl. Phys. Lett.*, vol. 85, no. 20, pp. 4795–4797, 2004.
- [82] E. J. Candès, X. Li, and M. Soltanolkotabi, "Phase retrieval via Wirtinger flow: Theory and algorithms," *IEEE Trans. Inf. Theory*, vol. 61, no. 4, pp. 1985–2007, Apr. 2015.
- [83] M. Mondelli and A. Montanari, "Fundamental limits of weak recovery with applications to phase retrieval," in *Proc. Conf. Learn. Theory*, 2018, pp. 1445–1450.
- [84] E. Wolf, "Three-dimensional structure determination of semi-transparent objects from holographic data," *Opt. Commun.*, vol. 1, no. 4, pp. 153–156, 1969.
- [85] A. Devaney, "Inverse-scattering theory within the Rytov approximation," *Opt. Lett.*, vol. 6, no. 8, pp. 374–376, 1981.
- [86] E. Soubies, T.-A. Pham, and M. Unser, "Efficient inversion of multiple-scattering model for optical diffraction tomography," *Opt. Exp.*, vol. 25, no. 18, pp. 21786–21800, 2017.
- [87] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proc. Int. Conf. Mach. Learn.*, 2017, pp. 214–223.
- [88] J. Latz, "On the well-posedness of Bayesian inverse problems," *SIAM/ASA J. Uncertainty Quantification*, vol. 8, no. 1, pp. 451–482, 2020.
- [89] A. Gelman, W. R. Gilks, and G. O. Roberts, "Weak convergence and optimal scaling of random walk metropolis algorithms," *Ann. Appl. Probability*, vol. 7, no. 1, pp. 110–120, 1997.
- [90] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. IEEE*, vol. 86, no. 11, pp. 2278–2324, Nov. 1998.
- [91] H. Xiao, K. Rasul, and R. Vollgraf, "Fashion-MNIST: A novel image dataset for benchmarking machine learning algorithms," 2017, *arXiv:1708.07747*.
- [92] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *Proc. Int. Conf. Learn. Representations*, 2015.
- [93] J. Lim, A. Goy, M. H. Shoreh, M. Unser, and D. Psaltis, "Learning tomography assessed using mie theory," *Phys. Rev. Appl.*, vol. 9, Mar. 2018, Art. no. 034027. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevApplied.9.034027>
- [94] U. S. Kamilov et al., "Learning approach to optical tomography," *Optica*, vol. 2, no. 6, pp. 517–522, 2015.
- [95] A. Durmus and É. Moulines, "On the geometric convergence for mala under verifiable conditions," 2022, *arXiv:2201.01951*.